

Analyzing the Effects of Comfort Relaxation on Energy Demand Flexibility of Buildings: A Multiobjective Optimization Approach

Pilar Morales-Valdés¹, Antonio Flores-Tlacuahuac¹, Victor M. Zavala^{2*}

¹ Departamento de Ingeniería y Ciencias Químicas, Universidad Iberoamericana,
Prolongación Paseo de la Reforma 880, México D.F., 01210, México

²Mathematics and Computer Science Division, Argonne National Laboratory,
9700 South Cass Avenue, Argonne IL 60439, United States

September 18, 2014

Abstract

We present a multiobjective optimization framework to evaluate the effects of comfort relaxation on the energy flexibility of buildings. This work is motivated by recent interest in understanding demand elasticity available for real-time electricity market operations and demand response events. We analyze the flexibility provided by an economics-based control architecture that directly minimizes total energy and by a traditional tracking control system that minimizes deviations from reference temperature and relative humidity set-points. Our study provides the following insights: (i) using percentage mean vote (PMV) and predicted percentage dissatisfied (PPD) constraints within an economics-based system consistently gives the most flexibility as comfort is relaxed, (ii) using PMV and PPD penalization objectives results in high comfort volatility, (iii) using temperature and relative humidity bounds severely overestimates flexibility, and (iv) tracking control offers limited flexibility even if used with optimal set-back conditions. We present a strategy to approximate nonlinear comfort regions using linear polyhedral regions, and we demonstrate that this reduces the computational complexity of optimal control formulations.

Keywords: multiobjective optimization, optimal control, HVAC systems, comfort relaxation, energy flexibility.

1 Introduction

Commercial buildings are valuable assets to power grid operators because they can enable demand elasticity [19]. Such flexibility is necessary to accommodate intermittent renewable power at a large scale and to avoid the construction of new generation plants. Demand flexibility can be achieved in buildings by shifting the demand profile in time and by relaxing comfort conditions to shed demand.

* Author to whom correspondence should be addressed. E-mail: vzavala@mcs.anl.gov, phone:+1 (630) 252-3343

Assessing the economic benefits that comfort relaxation can bring is nontrivial because system flexibility is a function of many factors such as the building design, real-time weather conditions, control architecture, and occupant acceptance [21, 20].

The heating, ventilation, and air conditioning (HVAC) system of a building comprises a large number of equipment units and material and energy resources that are monitored and coordinated in real time by a building management or control system (BMS). Thermal comfort and air quality conditions need to be enforced as occupancy levels, ambient conditions, and energy prices change. Two main control architectures are used in BMS systems. The traditional architecture (still dominating industry) determines set-points for equipment units such as delivery temperatures and economizer positions (recycle rates) as well as set-points for internal conditions such as zone temperatures, relative humidity, and pressure. Such set-points usually are tuned by experienced operators or by operational rules embedded in the BMS system. The set-points are then tracked by low-level, single-loop controllers such as thermostats. A key question that arises under this architecture is how to properly predict the amount of demand that the system will use for a given combination of set-points. This is particularly difficult because of the complex feedbacks and physical interactions that exist between equipment units and controllers. Limited knowledge of these interactions introduces a disconnect between global economic performance (e.g., total energy demand) and low-level control performance (e.g., set-point tracking) and can result in severe inefficiencies. This disconnect has been widely studied in the chemical industry [27, 25, 6].

Because of the inefficiencies of traditional control architectures, the building automation industry has recently shifted its interest to model-based management systems (also known as predictive control systems) [11, 5, 18, 22, 33]. These supervisory control architectures use dynamic models to predict the interactions between global HVAC variables and local zone conditions. In addition, they can directly optimize economic objective functions of different forms [7]. Consequently, these systems are also referred to as economics-based control systems [17]. All these features allow these advanced systems to predict and trade off energy (or cost) and comfort by manipulating multiple building variables simultaneously.

The vast majority of the building industry use neither comfort models nor occupant feedback routinely in operations. Consequently and, contrary to what is normally believed, most buildings operate under poor comfort conditions [14]. In addition, and to the best of our knowledge, only limited insights are available in the literature on the energy flexibility provided by different control systems as comfort conditions are relaxed. One can find many control formulations reported for which economic and energy savings potentials are evaluated. We refer the reader to the studies reported in [36, 7, 15, 30, 11, 5, 18, 22, 33, 31] and the comprehensive review [10]. None of these studies evaluates energy flexibility and control system behavior under relaxed comfort conditions.

The poor comfort performance of legacy control architectures results in resistance by occupants and building owners to consider emerging automation technologies. In addition, the limited understanding about the trade-offs between economic performance and comfort makes it difficult to fully appreciate the economic value of new technologies over prevailing ones and thus can make it difficult to commercialize them. We believe that performing more studies to evaluate these trade-offs is necessary to accelerate the deployment of new technologies. This, in turn, is essential to enabling

demand response and demand elasticity at a large scale [20, 35].

Several multiobjective optimization studies for buildings are available in the literature. In [24, 16] control studies are presented on the competing effects of air quality and energy consumption to demonstrate that significant energy reductions on energy use are possible with small relaxations of air quality conditions. In [9] the authors analyze the benefits of using multi-objective optimization in energy retrofit tasks. None of these works analyze trade-offs between energy use and thermal comfort.

In this work, we present a multiobjective optimization framework to evaluate the impact of comfort relaxation on the energy flexibility of different control systems. We compare the flexibility of different economics-based and traditional control architectures reported in the literature and used in practice. To perform our studies, we use a physical model of a single-zone building conditioned by an air-handling unit (AHU), and we incorporate a detailed PMV/PPD thermal comfort model. We emphasize that the intent of the multiobjective framework presented is not to obtain absolute numbers on the impact of comfort relaxation on demand flexibility. Such a study would require the consideration of many factors such as climate, building and HVAC design, and operational conditions. Instead, the intent of our study is to provide insights into the type of biased comfort perceptions and system volatility that can arise if inefficient control architectures and inappropriate comfort metrics are used in the control formulation. In addition, we seek to highlight the advantages that a multiobjective setting provides for analyzing and quantifying the benefits of economics-based control technologies.

The paper is structured as follows. In Section 2 we describe the dynamic model of the HVAC system and the thermal comfort model used in the optimal control formulations. In Section 3 we present the multiobjective framework used to analyze the behavior of different control formulations and comfort relaxation strategies. A numerical study is presented in Section 4, and computational issues are discussed in Section 5. Conclusions and future work are discussed in Section 6.

2 Dynamic Model of HVAC System

The building model considered in this work was presented in [28]. We use the thermal comfort model described in the ASHRAE standard 55-2004 [1]. The building model seeks to capture the effect of different global control variables on energy demand and local zone conditions. In particular, the model captures the dynamics of the zone CO₂ concentration, humidity, pressure, and temperature as well as the behavior of the AHU and the recycle temperatures, flows, and concentrations. We only describe the main variables of interest in the narrative. The full model notation and parameters are presented in Appendix A.

2.1 Material Balances

The total mass balance in the building zone is given by

$$\frac{dm(\tau)}{d\tau} = \rho \cdot (q^{in}(\tau) - q^{out}(\tau)) \quad (2.1)$$

where τ denotes time and $q^{in}(\cdot), q^{out}(\cdot)$ are inlet and outlet zone flows, respectively. The individual component dynamic mass balances are described by

$$V \cdot \frac{dC_i(\tau)}{d\tau} = q^{in}(\tau) \cdot C_i^{in}(\tau) - q^{out}(\tau) \cdot C_i(\tau) + n(\tau) \cdot n_{tot} \cdot G_i, \quad i \in \{CO_2, H_2O\}. \quad (2.2)$$

The occupancy signal of the space is given by $n(\tau)$, which takes a value of 1 if the space is occupied and a value of zero if it is unoccupied at a given time τ . The total number of occupants during occupied times is n_{tot} .

If we assume air with constant density and heat capacity the mass balance in the recycle is

$$q^{out}(\tau) + q^{amb}(\tau) = q^{ex}(\tau) + q^m(\tau) \quad (2.3a)$$

$$C_i(\tau) \cdot q^{out}(\tau) + C_i^{amb}(\tau) \cdot q^{amb}(\tau) = C_i(\tau) \cdot q^{ex}(\tau) + C_i^m(\tau) \cdot q^m(\tau), \quad i \in \{CO_2, H_2O\}. \quad (2.3b)$$

The mass balances in the AHU are given by,

$$q^{in}(\tau) = q^m(\tau) \quad (2.4a)$$

$$m_i^{rm}(\tau) = q^{in}(\tau) \cdot C_i^{in} - q^m(\tau) \cdot C_i^m(\tau). \quad (2.4b)$$

The mass removal rates in the AHU are denoted as m_i^{rm} . We consider $m_{CO_2}^{rm} = 0$ because this component is not removed in the AHU (only moisture is removed or added). The relationship between the zone pressure, mass, and temperature is given by the ideal gas law,

$$P(\tau) = \frac{m(\tau) \cdot R \cdot T(\tau)}{M \cdot V}. \quad (2.5)$$

The zone relative humidity is given by

$$RH(\tau) = 100 \cdot \frac{C_{H_2O}(\tau)}{C_{H_2O}^{sat}(\tau)}, \quad (2.6)$$

where $C_{H_2O}^{sat}(\tau)$ is the saturation density (concentration) and is computed from Antoine's equation [3],

$$\log_{10}(C_{H_2O}^{sat}(\tau)) = 8.07131 - \frac{1730.63}{T(\tau) - 39.73}. \quad (2.7)$$

The concentration of CO_2 in parts per million (ppm) is computed from

$$ppmV_{CO_2}(\tau) = 1000 \cdot \frac{C_{CO_2}(\tau) \cdot R \cdot T(\tau)}{M_{CO_2} \cdot P(\tau)}. \quad (2.8)$$

We note that interactions exist between different physical variables that are neither captured nor exploited by standard control architectures. For instance, relative humidity is affected by the saturation density, and this in turn is affected by the zone temperature. In addition, the CO_2 ppmV concentration is affected by the zone pressure and temperature. In [28] we noted that exploiting these variable interactions can give a control system high flexibility to minimize energy and satisfy air quality and comfort conditions. For instance, one can decrease the relative humidity by increasing the saturation density, not just by removing moisture in the AHU. In addition, one can decrease CO_2 concentration by increasing pressure because this implicitly increases the total mass of air in the zone. Standard architectures use single-loop controllers that cannot capture these global interactions. Instead, relative humidity and pressure conditions are usually kept fixed.

2.2 Energy Balances

The dynamic energy balance of the zone is given by

$$\frac{m(\tau)}{\rho} \cdot \frac{dT(\tau)}{d\tau} = q^{in}(\tau) \cdot T^{in}(\tau) - q^{out}(\tau) \cdot T(\tau) - \frac{U^w \cdot A^w}{\rho \cdot c_p} \cdot (T(\tau) - T^{amb}(\tau)) + \frac{n(\tau) \cdot n^{oc} \cdot Q^{oc}}{\rho \cdot c_p}. \quad (2.9)$$

Here, $T(\cdot)$ is the zone temperature, $T^{in}(\cdot)$ is the air supply temperature, and $T^{amb}(\cdot)$ is the ambient temperature. We consider only heat gains generated from the occupants. The energy balance in the recycle is

$$q^{out}(\tau) \cdot T(\tau) + q^{amb}(\tau) \cdot T_{amb}(\tau) = q^{ex}(\tau) \cdot T(\tau) + q^{in}(\tau) \cdot T^m(\tau). \quad (2.10)$$

Here, $T^m(\cdot)$ is the mixed stream temperature entering the AHU. Symbols $q^{amb}(\cdot)$, $q^{ex}(\cdot)$, and $q^{in}(\cdot)$ denote the ambient flow, exhaust flow, and delivery flow, respectively. The amount of moisture added/removed in the AHU $m_{H_2O}^{rm}(\cdot)$ is assumed to be proportional to the latent energy added/removed $Q^{lat}(\cdot)$. This is modeled as

$$Q^{lat}(\tau) = h^{lat} \cdot m_{H_2O}^{rm}(\tau). \quad (2.11)$$

The amount of sensible energy added/removed in the AHU is given by

$$Q^{sens}(\tau) = q^{in}(\tau) \cdot \rho \cdot c_p \cdot (T^{in}(\tau) - T^m(\tau)). \quad (2.12)$$

The total energy consumed by the HVAC system is computed from

$$Q^{hvac}(\tau) = |Q^{lat}(\tau)| + |Q^{sens}(\tau)|, \quad (2.13)$$

where $|\cdot|$ is the absolute value function. We do not consider energy consumed by fans. We reformulate equation (2.13) in differentiable form by using the following auxiliary equations:

$$Q^{hvac}(\tau) = (Q_+^{lat}(\tau) + Q_-^{lat}(\tau)) + (Q_+^{sens}(\tau) + Q_-^{sens}(\tau)) \quad (2.14a)$$

$$Q^{sens}(\tau) = Q_+^{sens}(\tau) - Q_-^{sens}(\tau) \quad (2.14b)$$

$$Q^{lat}(\tau) = Q_+^{lat}(\tau) - Q_-^{lat}(\tau). \quad (2.14c)$$

Here, $Q_+^{lat}(\tau)$, $Q_-^{lat}(\tau)$, $Q_+^{sens}(\tau)$, $Q_-^{sens}(\tau) \geq 0$ are auxiliary variables.

2.3 Thermal Comfort Model

Thermal comfort is the result of the heat exchange between the body and the environment, and it is influenced by four main variables (air temperature, radiant temperature, relative humidity, air speed) and two personal parameters (clothing and activity level or metabolic rate) [4, 1]. The comfort index used in the ASHRAE 55-2004 and the ISO 7730 standards [1, 13] is the predicted mean vote (PMV), originally proposed by Fanger [23]. We use the comfort model reported in Appendix D of the ASHRAE 55-2004 standard [1]. We have slightly adjusted its notation in order to enhance readability.

The PMV index reflects the thermal sensation during stationary conditions (approximately a settling time of an hour is assumed). The PMV is measured on a scale of seven values defined as follows: +3 hot, +2 warm, +1 slightly warm, 0 neutral, -1 slightly cool, -2 cool, -3 cold. A human is in a state of thermal equilibrium when the amount of heat produced and gained by the body is equal to amount of heat lost by the body. Fanger's model assumes that maximum comfort is achieved when such an equilibrium is reached. In the PMV scale, this corresponds to a value of zero. Acceptable comfort conditions are typically assumed to be in the PMV range $[-0.5, +0.5]$.

The PMV index can be calculated from the following heat balance relationship:

$$PMV(\tau) = (0.303 \cdot \exp(-0.036 \cdot M_m) + 0.028) \cdot L(\tau). \quad (2.15)$$

Here, $L(\tau)$ is the thermal load of the body and represents the difference between the body production and the loss of heat. $L(\tau)$ is calculated from the following contributions [1]:

$$L(\tau) = (M_m - M_w) - L_1(\tau) - L_2(\tau) - L_3(\tau) - L_4(\tau) - L_5(\tau) - L_6(\tau) \quad (2.16a)$$

$$L_1(\tau) = 3.05 \times 10^3 \cdot (5733 - 6.99 \cdot (M_m - M_w) - P_{H_2O}(\tau)) \quad (2.16b)$$

$$L_2(\tau) = 0.42 \cdot (M_m - M_w - 58.15) \quad (2.16c)$$

$$L_3(\tau) = 1.7 \times 10^{-5} \cdot M_m \cdot (5867.0 - P_{H_2O}(\tau)) \quad (2.16d)$$

$$L_4(\tau) = 1.4 \times 10^{-3} \cdot M_m \cdot (307.15 - T(\tau)) \quad (2.16e)$$

$$L_5(\tau) = 3.96 \times 10^{-8} \cdot f_{cl} \cdot (T^{cl}(\tau)^4 - T(\tau)^4) \quad (2.16f)$$

$$L_6(\tau) = f_{cl} \cdot h_c \cdot (T^{cl}(\tau) - T(\tau)). \quad (2.16g)$$

The different contributions to $L(\tau)$ (in order of appearance) correspond to net body heat production, heat losses through the skin, heat losses from sweating, latent heat losses from respiration, heat losses from dry respiration, heat losses from radiation, and heat losses by convection. Symbol $P_{H_2O}(\tau)$ denotes the water partial pressure in the air, $T^{cl}(\tau)$ is the clothing surface temperature, h_c is the convection heat transfer coefficient (a function of the relative air speed), I_{cl} is the clothing thermal insulation, f_{cl} is the clothing area factor, M_m is the metabolic heat gain, and M_w is the body external work. For simplicity in the presentation, we assume the mean radiant temperature to be equal to the zone temperature, and we assume a constant relative air speed.

The water partial pressure in the air $P_{H_2O}(\tau)$ (in Pascals) can be calculated from the saturation density $C_{H_2O}^{sat}(\tau)$ (in gr/m^3 or mmHg) and the relative humidity $RH(\tau)$ from [26, 2]:

$$P_{H_2O}(\tau) = \frac{133}{100} \cdot RH(\tau) \cdot C_{H_2O}^{sat}(\tau). \quad (2.17)$$

The surface temperature of the clothing $T^{cl}(\cdot)$ is obtained from the implicit algebraic equation

$$0 = a_4(\tau) + a_3(\tau) \cdot h_c - a_2 \cdot T^{cl}(\tau)^4 - T^{cl}(\tau) \cdot (1 + a_1 \cdot h_c), \quad (2.18)$$

where

$$a_1 = I_{cl} \cdot f_{cl} \quad (2.19a)$$

$$a_2 = 3.96 \times 10^{-8} \cdot a_1 \quad (2.19b)$$

$$a_3(\tau) = a_1 \cdot T(\tau) \quad (2.19c)$$

$$a_4(\tau) = 308.7 - 0.028 \cdot (M_m - M_w) + a_2 \cdot T(\tau)^4. \quad (2.19d)$$

The predicted percentage dissatisfied (PPD) is an index closely related to PMV. This represents the percentage of people that are expected not to be comfortable in a given thermal environment. This index is calculated by the following correlation:

$$PPD(\tau) = 100 - 95 \cdot \exp(-0.033 \cdot PMV(\tau)^4 - 0.218 \cdot PMV(\tau)^2) \quad (2.20)$$

The relationship between PPD and PMV is illustrated in Figure 1. Here, we highlight the points for PMV in the regions $[-0.5, +0.5]$ and $[-1, +1]$ and the corresponding PPD values, which are approximately 10% and 25%, respectively. According to this comfort model, the minimum possible value for PPD is 5%, which occurs at a PMV value of zero. In Figure 2 we visualize the comfort zone in a psychometric chart for values of PMV in the domain $[-0.5, +0.5]$ and $[-1, +1]$.




Figure 1: Relationship between PPD and PMV. The vertical lines indicate the $-0.5 \leq PMV \leq +0.5$ and $-1 \leq PMV \leq +1$ comfort regions. The horizontal lines indicate the corresponding PPD values, which are 10% and 25%, respectively, and the point of maximum comfort at 5%.

We note that Fanger's comfort model involves complex nonlinear equations that add significant difficulty when solved in conjunction with the material and energy balance equations. In Section 5 we present a strategy to approximate the comfort region predicted by Fanger's model using linear constraints and this ameliorates complexity.




Figure 2: Comfort zone visualization in a psychrometric chart. The light gray area corresponds to the $-0.5 \leq PMV \leq +0.5$ comfort region, and the dark gray area corresponds to the $-1 \leq PMV \leq +1$ region. Level curves are of constant relative humidity.

3 Multiobjective Control Formulation

Central to our study is the following multiobjective optimal control formulation,

$$\min_{u(\cdot)} \{ \Phi_e(x(\cdot), y(\cdot), u(\cdot)) , \Phi_c(x(\cdot), y(\cdot), u(\cdot)) \} \quad (3.21a)$$

$$\text{s.t.} \quad (3.21b)$$

$$\frac{dx}{d\tau} = f_x(x(\tau), y(\tau), u(\tau), d(\tau)), \tau \in [0, T] \quad (3.21c)$$

$$0 = f_y(x(\tau), y(\tau), u(\tau), d(\tau)), \tau \in [0, T] \quad (3.21d)$$

$$0 \leq h(x(\tau), y(\tau), u(\tau), d(\tau)), \tau \in [0, T] \quad (3.21e)$$

$$x(0) = \text{given.} \quad (3.21f)$$

Here, $x(\tau)$ are the dynamic state trajectories (temperature $T(\cdot)$, species concentrations $C_i(\tau)$, and air mass $m(\tau)$); $y(\tau)$ are the algebraic state trajectories (relative humidity $RH(\cdot)$, $PMV(\cdot)$, $PPD(\cdot)$); $u(\tau)$ are the control trajectories (sensitive heat load $Q^{sens}(\cdot)$, latent heat load $Q^{lat}(\cdot)$, ambient flow q^{amb} , and exhaust flow $q^{ex}(\cdot)$); and $d(\tau)$ are the system disturbance trajectories (ambient temperature $T^{amb}(\cdot)$, ambient concentrations $C_i^{amb}(\cdot)$, and occupancy signal $n(\cdot)$). The differential equations (3.21c) are the dynamic balances, while the algebraic equations (3.21d) comprise the rest of the model equations (balances and comfort equations). The inequality constraints (3.21e) denote bounds on air quality and flows. Comfort bounds are imposed on this control formulation separately.

Symbol $\Phi_e(\cdot)$, represents an economic objective function, while $\Phi_c(\cdot)$ represents a comfort objective. When these objectives are conflicting, we cannot simultaneously minimize them. Instead, the

optimal control solution is expressed in terms of the so-called Pareto front, which represents a set of solutions obtained by minimizing one objective by relaxing (or sacrificing) the other one. The Pareto front can be computed by using the so-called ϵ -constrained method. In this method, one of the objectives is moved into the constraints and is relaxed by a factor ϵ , which is a fixed parameter. This gives the following problem:

$$\min \quad \Phi_e(\cdot) \quad (3.22a)$$

$$\text{s.t.} \quad (3.21c) - (3.21e) \quad (3.22b)$$

$$\Phi_c(\cdot) \leq \epsilon. \quad (3.22c)$$

Here, we use the simplified notation $\Phi_e(\cdot) = \Phi_e(x(\cdot), y(\cdot), u(\cdot))$ and $\Phi_c(\cdot) = \Phi_c(x(\cdot), y(\cdot), u(\cdot))$. The relaxation factor ϵ is varied in the range $[\epsilon^L, \epsilon^U]$, where ϵ^U is obtained by solving the optimal control problem to minimize $\Phi_e(\cdot)$ and completely ignoring $\Phi_c(\cdot)$. The lower value ϵ^L is obtained by solving the optimal control by minimizing $\Phi_c(\cdot)$ and completely ignoring $\Phi_e(\cdot)$. The solution of problem (3.22) for a given relaxation factor ϵ gives a point along the Pareto front.

Any solution along the Pareto front also has a weight $\omega \in [0, 1]$ corresponding to the solution of the weighted-objective problem,

$$\min \quad (1 - \omega) \cdot \Phi_e(\cdot) + \omega \cdot \Phi_c(\cdot) \quad (3.23a)$$

$$\text{s.t.} \quad (3.21c) - (3.21e). \quad (3.23b)$$

The Pareto front can also be constructed by varying the weight ω on the range $[0, 1]$. When $\omega = 0$, the system minimizes the economic objective and comfort is ignored. When $\omega = 1$, the comfort objective is minimized, and economics are ignored. In other words, ω is the penalization weight of the comfort metric.

The economic objective function $\Phi_e(\cdot)$ used here is the accumulated HVAC energy demand over the time horizon $[0, N]$,

$$\Phi_e(\cdot) = \int_0^N Q^{hvac}(\tau) d\tau. \quad (3.24)$$

We use different metrics to define the comfort objective function $\Phi_c(\cdot)$. The first metric is the time-average PMV,

$$\Phi_c^{PMV,avg}(\cdot) = \frac{1}{N} \int_0^T PMV(\tau)^2 d\tau. \quad (3.25)$$

PMV has been squared because this metric can take negative and positive values and maximum comfort is obtained at $PMV=0$. This comfort penalization strategy has been used in [36, 8] using the weighted formulation (3.23). If the time-average PMV is the metric of comfort choice, we construct the Pareto front by solving problem (3.22) with relaxed $\Phi_c(\cdot)$. Because of the connection with the weighted objective strategy, this strategy is equivalent to minimizing the energy $\Phi_e(\cdot)$ while penalizing the time-average deviations from the point of maximum comfort ($PMV=0$). Because of this, we will refer to this strategy as *PMV penalization*.

Another way to enforce comfort is to impose constraints on PMV at each point in time $\tau \in [0, T]$. We can thus construct a Pareto front by progressively relaxing the PMV bounds. We do so by solving the following problem:

$$\min \quad \Phi_e(\cdot) \quad (3.26a)$$

$$\text{s.t.} \quad (3.21c) - (3.21e) \quad (3.26b)$$

$$-\epsilon \leq PMV(\tau) \leq \epsilon, \quad t \in [0, N]. \quad (3.26c)$$

Here, ϵ is varied in the PMV range $[-5, +5]$. We refer to this strategy as the *PMV constrained* strategy. This strategy has been used in [36, 12]. Problem (3.26) imposes the PMV bounds at each point in time, while (3.25) penalizes the time average. In other words, the *PMV penalization* strategy allows PMV bounds to be violated at certain times as long as the average remains bounded. Because of this, the *PMV objective* strategy is a relaxation of the *PMV constrained* strategy and will underestimate the energy demand. We demonstrate this in Section 5. One can show that the *PMV constrained* strategy is equivalent to constructing the Pareto front by solving problem (3.22) using the worst-case comfort metric,

$$\Phi_c^{PMV, max}(\cdot) = \max_{t \in [0, N]} |PMV(\tau)|. \quad (3.27)$$

The reason is that the minimum value of this metric is the minimum value of the problem

$$\min_{\epsilon} \quad \epsilon \quad \text{s.t.} \quad -\epsilon \leq PMV(\tau) \leq \epsilon, \quad t \in [0, N]. \quad (3.28)$$

The worst-case PMV metric is nonsmooth, so it is not amenable for use in the weighted-objective formulation (3.23). The worst-case metric is discussed only in order to achieve consistency between the different concepts presented here.

We will also study the time-average PPD comfort metric,

$$\Phi_c(\cdot) = \frac{1}{N} \int_0^T PPD(\tau) d\tau. \quad (3.29)$$

The PPD metric can take only positive values, so we can directly minimize its value. We construct the Pareto front by solving (3.22) and relaxing the time-average PPD metric (3.29). We refer to this strategy as the *PPD penalization* strategy. The commercial BuildingIQ system [33] follows this approach by using the weighted formulation (3.23). A practical advantage of the PMV and PPD penalization strategies is that they can lead to enhanced numerical robustness. In addition, they can be easily implemented by using black-box building and thermal comfort models. PPD is often preferred over PMV because it is easier to interpret. In Section 4 we show, however, that PPD induces more aggressive behavior.

We can also enable comfort relaxation by imposing bounds on PPD at each point in time and by progressively relaxing them. We do so by solving the following problem:

$$\min \quad \Phi_e(\cdot) \quad (3.30a)$$

$$\text{s.t.} \quad (3.21c) - (3.21e) \quad (3.30b)$$

$$PPD(\tau) \leq \epsilon, \quad t \in [0, N]. \quad (3.30c)$$

Here, ϵ is changed in the PPD range $[5, 100]$ where the lower bound is the minimum possible value for this metric obtained at a PMV value of zero. We refer to this strategy as the *PPD constrained* strategy. Problem (3.30) imposes the PPD bounds at each point in time, while (3.29) penalizes the time average. Consequently, (3.30) is a relaxation of (3.29) and will tend to underestimate the energy demand. The *PPD constrained* strategy is equivalent to constructing the Pareto front by relaxing the following worst-case PPD metric:

$$\Phi_c^{PPD,max}(\cdot) = \max_{t \in [0, N]} |PPD(\tau)|. \quad (3.31)$$

The third comfort metric explored here is the deviation from a given set-point. This is given by

$$\Phi_c(\cdot) = \frac{1}{N} \int_0^N (T(\tau) - T^c)^2 + (RH(\tau) - RH^c)^2 d\tau. \quad (3.32)$$

Here, T^c and RH^c are temperature and relative humidity set-points, respectively. This comfort metric will be key in exploring the performance of different control strategies widely used in industry and academic studies. This metric was used to evaluate the impact of comfort relaxation on energy demand in [28, 34].

The first strategy consists of a set-point tracking control architecture that minimizes (3.32) in order to move as fast as possible to a given set-point. Typically, this set-point is set to a point of maximum comfort (PMV=0). This typically occurs around a temperature of 22 °C and a relative humidity of 50%. Consequently, solving problem (3.22) by relaxing this metric is equivalent to allowing the set-point tracking controller to deviate from the point of maximum comfort. This strategy was considered in [36] where the authors designed a predictive controller that penalizes deviations in temperature and relative humidity from a point of maximum comfort. We refer to this strategy as the *maximum comfort tracking* strategy. This strategy is used in practice because it avoids the need of incorporating the nonlinear thermal comfort model in the optimal control formulation, thus avoiding its computational complexity. In addition, this strategy mimics the performance of typical PID control loops that are tuned to closely track set-points. Relaxing the comfort objective (3.32) can be seen as a *detuning* of the tracking controllers that allows them to reach the set-point less aggressively. As we show in Section 4, this strategy can result in significant distortions on energy demand and high volatility as comfort is relaxed.

Another strategy typically used in practice is the following. Because a set-point tracking control architecture is typically in place, the amount of energy consumed can be relaxed by changing the set-point values T^c , RH^c to values of relaxed comfort. These set-points are also known as set-back conditions. For instance, we might think of moving the temperature set-point from 23 °C to 25 °C to relax the PPD from 5% to 11%. We will refer to this strategy as *set-back tracking*. This strategy is widely used in commercial and residential buildings because it is simple and intuitive.

The last strategy considered and used in several studies consists of enforcing comfort by using

temperature and relative humidity bounds [5, 32, 22]. In this case, the following problem is solved:

$$\min \quad \Phi_e(\cdot) \quad (3.33a)$$

$$\text{s.t.} \quad (3.21c) - (3.21e) \quad (3.33b)$$

$$\epsilon_T^L \leq T(\tau) \leq \epsilon_T^U \quad (3.33c)$$

$$\epsilon_{RH}^L \leq RH(\tau) \leq \epsilon_{RH}^U. \quad (3.33d)$$

Here, $\epsilon_T^L, \epsilon_T^U$ are lower and upper bounds for the zone temperature, and $\epsilon_{RH}^L, \epsilon_{RH}^U$ are lower and upper bounds for relative humidity. This strategy is practical because it avoids the incorporation of the nonlinear thermal comfort model. Unfortunately, this strategy overestimates the feasible region imposed by PMV, as we show in Section 5. This overestimation is depicted in Figure 3 where we present the feasible PMV region for the range $[-0.5, 0.5]$ (light gray area) and the region imposed by the tightest upper and lower bounds for temperature and relative humidity (dark gray area). The overestimation can be significant. The overestimation of the feasible PMV region results in the underestimation of the energy demand, and this can lead to comfort violations. Because of this, we refer to this strategy as *comfort overrelaxation*.

We can construct a Pareto front for this strategy by finding the temperature and relative humidity bounds corresponding to a given PMV region. Most control studies use this strategy, but they only use the flexibility of the zone temperature to relax comfort [5, 18, 22, 17, 11, 19]. This implicitly assumes that the relative humidity is fixed, because it is rarely modeled. In this case, imposing bounds on temperature is equivalent to imposing bounds on PMV or PPD because one can always find the appropriate temperature bounds (at a fixed relative humidity) that correspond to the PMV bounds. Several studies have demonstrated that allowing relaxation of relative humidity can enable significant energy demand reductions [36, 28]. In Section 5 we use the multiobjective framework to demonstrate that this is indeed the case. In addition, we show that relaxing relative humidity also enhances flexibility.

4 Numerical Study

We present a numerical study to explore the energy flexibility of the different control strategies. To do so, we solve the different multiobjective optimal control formulations using a time horizon of 48 hours. The horizon is discretized in time intervals of one hour. We use real ambient temperature and relative humidity conditions at a location in the U.S. Midwest during a summer day [28]. The profiles are presented in Figure 4. We assume an ambient concentration of CO₂ of 400 ppmV, and we consider a building zone with a volume of 1,000 m³ and 500 occupants during occupied hours. We impose an upper bound on CO₂ concentration of 1,000 ppmV. All the simulation models needed to reproduce the results can be obtained in <http://www.mcs.anl.gov/~vzavala/MultiobjComfort.tgz>.

4.1 PMV Constrained vs. PMV Penalization

The Pareto fronts obtained with the *PMV constrained* and the *PMV penalization* strategies are presented in Figure 5. In the graph of the left panel, the vertical axis is the total amount of energy




Figure 3: Overestimation of the PMV feasible region using temperature and relative humidity bounds. The light gray area denotes a PMV region in the range $[-0.5, +0.5]$ and the dark gray area denotes the feasible region described by the tightest temperature and relative humidity bounds.




Figure 4: Ambient temperature (top), occupancy (middle), and relative humidity profiles (bottom).

consumed over the time horizon $\Phi_e(\cdot)$ while the horizontal axis is the time-average PPD metric $\Phi_c^{PPD,avg}(\cdot)$. We use time-average PPD to visualize the results, instead of the time-average PMV metric $\Phi_c^{PMV,avg}(\cdot)$, because it is easier to make inferences on comfort impact. We observe that the PMV penalization strategy underestimates the energy demand compared with that of the PMV con-

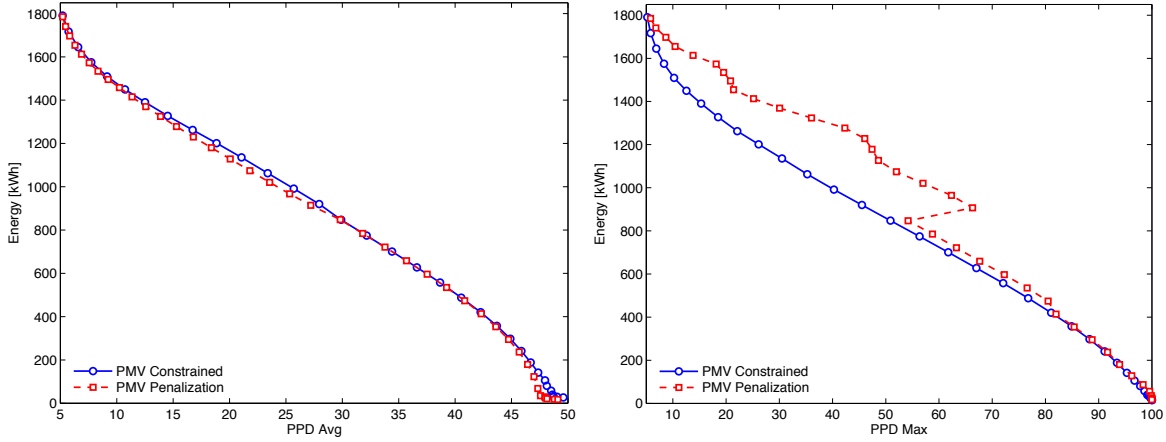


Figure 5: Pareto front for *PMV constrained* and *PMV penalization* strategies. Visualization using time-average PPD (left) and maximum PPD (right).

strained strategy. This result illustrates that penalization gives more flexibility to the control system. Thus, one might be tempted to conclude that it is a more efficient strategy. We now illustrate that this flexibility comes at the expense of significant volatility in the comfort profiles.

In the right panel of Figure 5 we visualize the Pareto fronts using the maximum PPD metric ($\Phi_c^{PPD,max}(\cdot)$) in the horizontal axis. The PMV penalization strategy gives inconsistent results. As comfort is relaxed, energy demand is reduced, as expected; the reduction, however, is nonmonotone. In addition, the departure of the *PMV penalization* front from the *PMV constrained* front indicates that the penalization strategy consumes more energy for the same level of the worst-case PPD. The gap can be as large as 300 kWh. Note also that at a value of maximum PPD of around 70%, the maximum PPD jumps back to a value of around 50%. In other words, almost identical energy demands are obtained for very different PPD values. After this break point, both fronts converge, which implies that PPD becomes constant in time (because average and worst-case values are the same). This erratic behavior of the *PMV penalization* strategy illustrates the difficulties encountered in controlling comfort profiles as conditions are relaxed (i.e., by penalizing the PMV term less and less). The *PMV constrained* strategy, on the other hand, gives a smooth transition between different comfort relaxation levels.

4.2 PPD Constrained vs. PPD Penalization

We now construct Pareto fronts for the *PPD constrained* and the *PPD penalization* strategies. The results are presented in Figure 6. From the left panel we can see that the *PPD penalization* strategy underestimates the amount of energy used. The effects, however, are much more pronounced in this case compared with *PMV penalization*. This can be attributed to the additional nonlinearity induced by the exponential PPD function (2.20). In the right panel of Figure 6 we can see that *PPD penalization* yields extremely high volatility in PPD profiles. At around a maximum PPD value of 20%, the Pareto front makes a sudden jump to a maximum value of 100%. In other words, while average PPD can be maintained at reasonable levels, a slight relaxation in average PPD can yield extremely high




Figure 6: Pareto front for *PPD constrained* and *PPD penalization* strategies. Visualization using time-average PPD (left) and maximum PPD (right).




Figure 7: Pareto front for *PMV constrained* and *PPD penalization* strategies.

discomfort levels at certain times. We thus conclude that while *PPD penalization* gives more energy shedding flexibility, this is not a robust strategy to induce comfort relaxation. From this result we also conclude that *PPD penalization* yields much worse volatility compared with *PMV penalization*. The *PPD constrained* strategy, on the other hand, exhibits smooth transitions between different comfort levels.

4.3 PMV Constrained vs. PPD Constrained

In Figure 7 we compare the *PMV constrained* and the *PPD constrained* strategies. The Pareto fronts are nearly identical. The reason is the nature of the PPD function, which imposes symmetry on both sides of PMV. We thus conclude that both strategies are equally effective at relaxing comfort. This conclusion is important because it is preferable to impose bounds in the space of PMV in order to avoid the nonlinear nature of the PPD function (2.20).




Figure 8: Pareto fronts for *economics-based* strategy with PMV constraints and for the *maximum set-point tracking* strategy.

4.4 Economic-Based vs. Maximum Comfort Tracking

We now compare the energy shedding potential of the *PMV constrained* strategy (which is an economics-based control strategy) against that of the *maximum comfort tracking* strategy. The results are presented in Figure 8. At maximum comfort (PPD of 5%) the total amount of energy consumed by the *PMV constrained* strategy is 1,800 kWh. A relaxation of the maximum PPD value to 10% reduces the energy level of 1,500 kWh. This is a shedding potential of 300 kWh or 16% in relative terms. A relaxation to a maximum PPD value of 20% yields a shedding potential of 500 kWh or 27%. At a maximum PPD value of 30% the potential is 700 kWh or 38%. The tracking strategy consistently uses more energy than does the economics-based strategy. At the maximum PPD level of 5% the relative difference in energy demand is of 10% (absolute difference of 200 kWh), while at a PPD level of 20% reaches a level of 30% (600 kWh). The increasing relative difference indicates that the economic-based strategy offers much more shedding flexibility as comfort is relaxed. For instance, relaxing the maximum PPD level for the economics-based strategy from a level of 5% to 25% offers a shedding potential of 600 kWh (1,800 kWh to 1,200 kWh). The tracking strategy offers a shedding potential of 300 kWh (2,000 kWh to 1,700 kWh). This is half of the flexibility of the economics-based strategy. The tracking strategy also exhibits high comfort volatility as conditions are relaxed. This result illustrates that tracking strategies are difficult to tune to give appropriate comfort relaxation.

4.5 Economic-Based vs. Set-Back Relaxation

We now compare the *economic-based* PMV constrained strategy against the *set-back relaxation* strategy. We fix the relative humidity set-point RH^c to a value of 50% and find the temperatures T^c at which PMV values on the range $[0,3]$ are obtained. This correspond to temperature set-points in the range $[26,33]^\circ\text{C}$. We use positive values of PMV because this is a cooling season. We then use the temperature set-points to form the Pareto front by minimizing $\Phi_c^{track}(\cdot)$ (3.32). To perform a fair comparison,




Figure 9: Pareto front for *economic-based PMV constrained* with fixed relative humidity, *economic-based PMV constrained* with free relative humidity, and *set-back* relaxation strategies.

we also fix the relative humidity to a value of 50% for the economic-based strategy. This fixed relative humidity strategy has been considered in [36]. To evaluate the additional flexibility that relative humidity offers, we also compare the performance of the economic-based strategy using free relative humidity. The results are presented in Figure 9. The set-back relaxation strategy indeed provides energy flexibility. In addition, its behavior is stable. This flexibility, however, is very limited compared with that of the economics-based strategy. In particular, the flexibility of the set-back strategy when relaxing comfort from a maximum PPD value of 5% to a value of 25% is less than 10% (200 kWh). The flexibility for the economic-based strategy using the same relaxation values is 20% (400 kWh). We can conclude that a tracking control architecture provides limited flexibility even if set-points are fully relaxed. The reason is that a tracking controller will always try to stay close to the set-point given and this will come at the expense of increased energy demand and limited flexibility.

By comparing the Pareto front for the economic-based strategy with and without fixed relative humidity we can see that optimizing for relative humidity yields significant reductions in energy levels and significant improvements in flexibility. Relaxing the maximum PPD from 5% to 25% results in an energy flexibility of 400 kWh for the fixed relative humidity strategy. For the free relative humidity strategy the flexibility is 600 kWh. This is an improvement of 50%.

We emphasize that the results presented here are to understand general trends of behavior. It would be interesting to perform an actual deployment to evaluate the trade offs considered here using measurement and verification techniques. Such deployment should consider the relaxation of relative humidity set-points simultaneously with those of temperature humidity to verify energy savings in the air-handling unit. Comfort should be verified using occupant polls to verify how close are PMV predictions to actual perceptions of comfort.

5 Computational Issues

In this section we demonstrate that one can approximate the comfort region described by the non-linear thermal comfort model by using linear constraints. This ability is important for reducing the complexity of the control formulation. We also compare the energy and comfort performance of our approximation with that of the *overrelaxation* strategy.

5.1 Energy and Comfort Performance

Consider the comfort region of Figure 3. The four corners of the comfort region can be obtained by fixing the relative humidity to minimum and maximum values allowable (in this case we use $RH_1=0\%$ and $RH_2=100\%$ but this can be set to more conservative values based on operational considerations). As an example, consider the PMV region $[-0.5,0.5]$ highlighted in Figure 3. For the relative humidity values RH_1 and RH_2 first consider a PMV value of -0.5 (left side). The temperatures are given by $T_1^{left} = 21$ and $T_2^{left} = 23^\circ\text{C}$. For the right side, we find the temperatures giving a PMV value of $+0.5$ at the given relative humidities. The values are $T_1^{right} = 24$ and $T_2^{right} = 27^\circ\text{C}$. The comfort region is then approximated using the polyhedral region described by the intersection of the following inequality constraints,

$$RH(\tau) \geq RH_1, \quad t \in [0, N] \quad (5.34a)$$

$$RH(\tau) \leq RH_2, \quad t \in [0, N] \quad (5.34b)$$

$$RH(\tau) \geq RH_2 + \frac{RH_2 - RH_1}{T_1^{left} - T_2^{left}} (T(\tau) - T_1^{left}), \quad t \in [0, N] \quad (5.34c)$$

$$RH(\tau) \leq RH_2 + \frac{RH_2 - RH_1}{T_1^{right} - T_2^{right}} (T(\tau) - T_1^{right}), \quad t \in [0, N]. \quad (5.34d)$$

For the PMV region $[-0.5,+0.5]$, the *overrelaxation strategy* imposes the relative humidity bounds $RH_1 \leq RH(\tau) \leq RH_2$ and the temperature bounds $T_1^{left} \leq T(\tau) \leq T_2^{right}$.

We compare the energy performance of this *polyhedral approximation* strategy against the performance of the temperature and relative humidity bounds formulation (3.33) (*overrelaxation strategy*) and against the performance of the *PMV constrained* formulation (3.26) for different PMV regions. We compute the resulting maximum PPD value for all the formulations to determine the amount of discomfort induced by the approximations. The results are presented in Table 1. For the $[-1,+1]$ region, the overrelaxation strategy underestimates the energy demand by 365 kWh. This represents 30% less energy compared with the *PMV constrained* strategy. The maximum PPD, however, reaches 60% while the *PMV constrained* strategy reaches a maximum of 26% (consistent with a PMV in $[-1,+1]$). In other words, 34% more people become dissatisfied under the overrelaxation strategy. For the $[-0.5,+0.5]$ region, 36% more occupants are dissatisfied. For the $[-1,+1]$ region, the polyhedral approximation strategy underestimates energy demand by 20 kWh. Only 3% more occupants become dissatisfied relative to the *constrained* strategy. The performance is dramatically improved and this trend remains if the comfort region is either relaxed or constrained.

We have not explicitly computed a Pareto for the *comfort overrelaxation* strategy but the results in Table 1 present three points along the front. In other words, we used the the temperature and relative

Table 1: Performance of different comfort bounding strategies.

PMV Region	Strategy	$\Phi_e(\cdot)$ [kWh]	$\Phi_c^{PPD,max}(\cdot)$ [%]	NLP Iterations
[-0.5,+0.5]	PMV	1,495	10.2	76
	Overrelaxation	1,222	46.8	23
	Polyhedral	1,493	10.8	15
[-1.0,+1.0]	PMV	1,201	26.1	57
	Overrelaxation	836	60.1	50
	Polyhedral	1,181	27.9	28
[-2.0,+2.0]	PMV	469	76.1	51
	Overrelaxation	161	95.9	11
	Polyhedral	429	79.3	32

humidity bounds corresponding to the increasingly relaxed PMV regions. As we have discussed, the overrelaxation strategy consistently underestimates energy demand compared to the *PMV constrained* strategy. In addition, it can be concluded that this strategy yields more volatility in comfort conditions as the bounds are relaxed.

5.2 Computational Performance

To solve the optimal control problems we use an implicit Euler discretization strategy and solve the resulting nonlinear programming (NLP) problems using IPOPT [29]. In Table 1 we present the number of NLP iterations required for the different formulations. The incorporation of Fanger’s thermal model equations and the PMV constraints introduce significant nonlinearities. The optimization solver requires 75 iterations in the worst case and 51 in the best case to converge. The *overrelaxation* strategy requires 50 iterations in the worst case and 11 in the best case. The *polyhedral approximation* strategy requires 32 iterations in the works case and 15 in the best case. The improved performance obtained with the *polyhedral approximation* over the PMV constrained strategy can be attributed directly to the elimination of the nonlinear constraints by linear constraints. The improvement over the *overrelaxation* strategy can be attributed to the fact that the polyhedral constraints narrow down the feasible region, thus reducing the search space of the solver.

6 Conclusions and Future Work

We have presented a multiobjective framework to analyze the effect of comfort relaxation on the energy flexibility of building controls systems. We analyzed the flexibility of two main paradigms: economic-based control strategies and tracking control strategies. We considered the use of constraints and the use of penalization terms on percentage mean vote and predicted percentage dissatisfied metrics. We have found that PPD and PMV penalization terms yield poor control of comfort conditions and high system volatility as comfort requirements are relaxed. Consequently, the use of direct PMV and PPD constraints is advised. We have also found that a tracking control architecture

offers limited energy flexibility as comfort requirements are relaxed. In addition, we have found that avoiding the use of detailed thermal comfort models in control formulations overestimates flexibility and leads to poor comfort conditions. We have proposed a strategy that approximates the comfort region described by Fanger’s thermal comfort model using linear constraints. This significantly reduces the computational complexity of economics-based control formulations. [The results presented here can have important consequences in actual applications. In particular, they point to the need of considering *co-optimization of temperature and relative humidity set-points*, which is rarely considered in practice. In addition, the results suggest that rigorous comfort models can be approximated reasonably using linear polyhedral regions and this can facilitate the implementation of optimization-based control techniques.](#)

A Model Nomenclature

A.1 Variables

$m(\cdot)$	zone total air mass, gr_{air}
$C_i(\cdot)$	zone concentration of component i , gr_i/m^3
$C_i^{in}(\cdot)$	concentration of component i in inlet air, gr_i/m^3
$C_i^m(\cdot)$	concentration of component i in mixer, gr_i/m^3
$C_{H_2O}^{sat}(\cdot)$	zone saturation density, gr_i/m^3
$m_i^{rm}(\cdot)$	mass removal rate of component i in AHU, gr_i/hr
$P(\cdot)$	zone pressure, atm
$ppmV_{CO_2}(\cdot)$	zone CO_2 concentration, $ppmV$
$PMV(\cdot)$	percentage mean vote, $-$
$PPD(\cdot)$	predicted percentage dissatisfied, $\%$
$q^{in}(\cdot)$	volumetric inlet flow rate, m^3/hr
$q^{out}(\cdot)$	volumetric outlet flow rate, m^3/hr
$q^m(\cdot)$	volumetric flow rate in mixed, m^3/hr
$q^{amb}(\cdot)$	volumetric ambient air flow rate, m^3/hr
$q^{ex}(\cdot)$	volumetric exhaust air flow rate, m^3/hr
$Q^{lat}(\cdot)$	latent heat removed in AHU, kJ/hr
$Q^{sens}(\cdot)$	sensible heat removed in AHU, kJ/hr
$Q^{hvac}(\cdot)$	total HVAC energy, kJ/hr
$P_{H_2O}(\cdot)$	water partial pressure, Pa
$RH(\cdot)$	zone relative humidity, $\%$
$T(\cdot)$	zone temperature, K
$T^{in}(\cdot)$	temperature of inlet air, K
$T^m(\cdot)$	temperature in mixer, K
$T^{cl}(\cdot)$	clothing surface temperature, K

A.2 Parameters

$C_i^{amb}(\cdot)$	concentration of component i in ambient air, gr/m^3
$T^{amb}(\cdot)$	temperature of ambient air, K
n^{oc}	number of occupants under occupied periods, 500
G_i^z	generation rate of component i per occupant, ($CO_2 = 2.4, H_2O = 50$) gr/hr
Q^{oc}	building heat gain per occupant, 432 kJ/hr
U^w	wall heat-transfer coefficient, $18\text{ kJ/hr} \cdot m^2 \cdot K$
V^z	total building volume, 1000 m^3
A^w	total wall heat-transfer area, 600 m^2
c_p	air heat capacity at standard conditions, $1.0 \times 10^{-3}\text{ kJ/gr} \cdot K$
h^{lat}	latent heat of condensation, 2.46 kJ/gr
ρ	air density at standard conditions, 1200 g/m^3
I_{cl}	insulation factor, 0.155
f_{cl}	clothing area factor, 1.15
M	air molecular weight, 29 gr/gr_{mol}
M_{CO_2}	CO_2 molecular weight, 44 gr/gr_{mol}
M_m	metabolic heat gain, 58.5 W/m^2
M_w	body external work, 0 W/m^2
R	universal gas constant, $0.082 \times 10^{-3}\text{ atm} \cdot m^3/\text{gr}_{mol}K$

Acknowledgments

This work was supported by the U.S. Department of Energy, under Contract No. DE-AC02-06CH11357.

References

- [1] ASHRAE Standard 55-2004. Thermal environment conditions for human occupancy. *American Society of Heating Ventilating and Air-conditioning Engineers*, pages 2–24, 2004.
- [2] Michaelsen B. and Eiden J. Human comfort modelica-library thermal comfort in buildings and mobile applications. *The Modelica Association*, pages 403–412, 2009.
- [3] Tashtoush B., Molhim M., and Al-Rousan M. Dynamic model of an hvac system for control analysis. *ASHRAE Transactions*, 104(1):145–167, 1998.
- [4] Olesen B. W. and Brager G. S. A better way to predict comfort. *ASHRAE Journal*, pages 20–26, 2004.
- [5] Sorin C Bengea, Anthony D Kelman, Francesco Borrelli, Russell Taylor, and Satish Narayanan. Implementation of model predictive control for an hvac system in a mid-size commercial building. *HVAC&R Research*, 20(1):121–135, 2014.

- [6] Lorenz T Biegler and Victor M Zavala. Large-scale nonlinear programming using ipopt: An integrating framework for enterprise-wide dynamic optimization. *Computers & Chemical Engineering*, 33(3):575–582, 2009.
- [7] J. E. Braun. Reducing energy costs and peak electricity demand through optimal control of building thermal storage. *ASHRAE Transactions*, 96(2):876–888, 1990.
- [8] M Castilla, JD Álvarez, M Berenguel, F Rodríguez, JL Guzmán, and M Pérez. A comparison of thermal comfort predictive control strategies. *Energy and buildings*, 43(10):2737–2746, 2011.
- [9] Christina Diakaki, Evangelos Grigoroudis, and Dionyssia Kolokotsa. Towards a multi-objective optimization approach for improving energy efficiency in buildings. *Energy and Buildings*, 40(9):1747 – 1754, 2008.
- [10] Anastasios I Dounis and Christos Caraiscos. Advanced control systems engineering for energy and comfort management in a building environment a review. *Renewable and Sustainable Energy Reviews*, 13(6):1246–1261, 2009.
- [11] Gregor P. Henze, Clemens Felsmann, and Gottfried Knabe. Evaluation of optimal control for active and passive building thermal storage. *International Journal of Thermal Sciences*, 43(2):173 – 183, 2004.
- [12] Jung-Ho Huh and Michael J Brandemuehl. Optimization of air-conditioning system operating strategies for hot and humid climates. *Energy and Buildings*, 40(7):1202–1213, 2008.
- [13] ISO-7730. Ergonomics of the thermal environment : Analytical determination and interpretation of thermal comfort using calculation of the pmv and ppd indices and local thermal comfort criteria. *Organisation for Standardisation*, pages 1–52, 2005.
- [14] P. Stluka K. Marik, J. Rojcek and J. Vass. Advanced HVAC control: Theory vs. reality. *Proceedings of 18th IFAC World Congress*, pages 3108–3113, 2011.
- [15] Fong K.F., Hanby V.I., and Chowa T.T. Hvac system optimization for energy management by evolutionary programming. *Energy and Buildings*, 38:220–231, 2006.
- [16] Hongbin Liu, SeungChul Lee, MinJeong Kim, Honglan Shi, Jeong Tai Kim, Kailas L. Wasewar, and ChangKyoo Yoo. Multi-objective optimization of indoor air quality control and energy consumption minimization in a subway ventilation system. *Energy and Buildings*, 66(0):553 – 561, 2013.
- [17] Jingran Ma, S Joe Qin, Bo Li, and Tim Salsbury. *Economic model predictive control for building energy systems*. IEEE, 2011.
- [18] Yudong Ma, F. Borrelli, B. Hancey, A. Packard, and S. Bortoff. Model predictive control of thermal energy storage in building cooling systems. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 392 –397, dec. 2009.

- [19] Mehdi Maasoumy, Catherine Rosenberg, Alberto Sangiovanni-Vincentelli, and D Callaway. Model predictive control approach to online computation of demand-side flexibility of commercial buildings hvac systems for supply following. In *Submitted to the American Control Conference (ACC)*, 2014.
- [20] Johanna L Mathieu, Phillip N Price, Sila Kiliccote, and Mary Ann Piette. Quantifying changes in building electricity use, with application to demand response. *Smart Grid, IEEE Transactions on*, 2(3):507–518, 2011.
- [21] Naoya Motegi, Mary Ann Piette, David S Watson, Sila Kiliccote, and Peng Xu. Introduction to commercial building control strategies and techniques for demand response. *Lawrence Berkeley National Laboratory LBNL-59975*, 2007.
- [22] F. Oldewurtel, A. Parisio, C.N. Jones, M. Morari, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and K. Wirth. Energy efficient building climate control using stochastic model predictive control and weather predictions. In *American Control Conference (ACC), 2010*, pages 5100–5105, 30 2010-july 2 2010.
- [23] Fanger P. O. Assessment of man’s thermal comfort in practice. *British Journal of Industrial Medicine*, 30:313–324, 1973.
- [24] Adams Rackes and Michael S. Waring. Using multiobjective optimizations to discover dynamic building ventilation strategies that can improve indoor air quality and reduce energy use. *Energy and Buildings*, 75(0):272 – 280, 2014.
- [25] James B Rawlings and Rishi Amrit. Optimizing process economic performance using model predictive control. In *Nonlinear Model Predictive Control*, pages 119–138. Springer, 2009.
- [26] Bartkevičius S., Račkienė R., and Virbalis J.A. Analysis of the energy balance in the system human- clothing- environment. *Electronics and electrical engineering*, 7(87):61–64, 2008.
- [27] A. Toumi, M. Diehl, S. Engell, H. G. Bock, and J. P. Schlöder. Finite horizon optimizing control of advanced smb chromatographic processes. *16th IFAC World Congress, Prague*, 2005.
- [28] Zavala V.M. Real-time optimization strategies for building systems. *Industrial and Engineering Chemistry Research*, 52(9):3137–3150, 2013.
- [29] A. Wächter and L. T. Biegler. On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106:25–57, 2006.
- [30] J. Wang, M. Shahidehpour, and Z. Li. Security-constrained unit commitment with volatile wind power generation. *IEEE Transactions on Power Systems*, 23:1319–1327, 2008.
- [31] V. M. Zavala, E. M. Constantinescu, M. Anitescu, and T. Krause. On-line economic optimization of energy systems using weather forecast information. *Journal of Process Control*, 19:1725–1736, 2009.

- [32] V. M. Zavala, E. M. Constantinescu, M. Anitescu, and T. Krause. On-line economic optimization of energy systems using weather forecast information. *Journal of Process Control*, 19:1725–1736, 2009.
- [33] V. M. Zavala, D. Skow, T. Celinski, and P. Dickinson. Techno-economic evaluation of a next-generation building energy management system. *Technical Report ANL/MCS-TM-313, Argonne National Laboratory*, 2011.
- [34] Victor M Zavala. Real-time resolution of conflicting objectives in building energy management: an utopia-tracking approach. *SimBuild*, 2012.
- [35] VM Zavala, C Thomas, M Zimmerman, and A Ott. Next-generation building energy management systems and implications for electricity markets. *technical Report ANL/MCS-TM315, Argonne National Laboratory (ANL)*, 2011.
- [36] Freire Z.R., Oliveira G. H., and Mendes N. Predictive controllers for thermal comfort optimization and energy savings. *ABCM Symposium series in mechatronics*, 3:839–848, 2008.

<p>The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.</p>
